

ANKAphase User Guide

Timm Weitkamp^{1,a}, David Haas^{2,3,b}, and Alexander Rack^{4,c}

¹Synchrotron Soleil, Gif-sur-Yvette, France

²ANKA Light Source, Karlsruhe Institute of Technology (KIT), Germany

³Hochschule Darmstadt, Germany

⁴European Synchrotron Radiation Facility (ESRF), Grenoble, France

^aE-mail: timm.weitkamp@synchrotron-soleil.fr

^bE-mail: david.haas@kit.edu

^cE-mail: arack@snafu.de

June 20, 2011

Citation and further reading: T. Weitkamp, D. Haas, D. Wegrzynek, and A. Rack, *ANKAphase: software for single-distance phase-retrieval from inline X-ray phase contrast radiographs*, Journal of Synchrotron Radiation (2011), vol. 18, issue 4, p. 617-629, DOI: 10.1107/S0909049511002895.

Contents

1	General description	1
1.1	What ANKAphase is	1
1.2	The phase reconstruction algorithm	1
1.3	Flat-field and dark-image correction	2
1.4	Boundary treatment; image padding	2
2	Where to get ANKAphase	3
3	Files in the distribution	3
4	Requirements	4
4.1	Platforms and operating systems	4
4.2	Software requirements	4
4.3	Data format requirements	5
5	Getting started	5
5.1	Running it as a standalone application	5
5.2	Running it as a plugin for ImageJ	5
5.3	Using it	6
5.3.1	General settings	6
5.3.2	Flat-field settings	7
5.3.3	Phase-retrieval settings	9

5.3.4	Other settings	10
5.3.5	Launching the processing	10
5.3.6	Menu bar	11
6	Further information; referencing	11

1 General description

1.1 What ANKAphase is

ANKAphase [1] is a computer program that processes X-ray inline phase-contrast radiographs and reconstructs the projected thickness of the object(s) imaged in the radiographs. It uses an algorithm described in a paper by David Paganin et al. [2] (see also section 1.2, below).

ANKAphase is designed to be used with sets of radiographs recorded for tomographic reconstruction (although it does not perform tomographic reconstruction itself). This means that it can process a series of phase-contrast radiographs taken of different views of a sample with an otherwise unchanged setup. If desired, it also performs flat-field and dark-image correction (i.e., normalization to the background intensity profile of the X-ray beam and subtraction of the “dark-signal” offset value for each pixel).

ANKAphase is written in Java and can either be used as a standalone application or as a plugin to ImageJ, a widely-used scientific image viewing and processing program developed and maintained at the National Institutes of Health (NIH), Bethesda, Maryland, USA (see sections 4 and 5.2 below).

1.2 The phase reconstruction algorithm

In this section we give a short synopsis of the algorithm used by *ANKAphase* and first published by Paganin et al. [2]. Although very robust under conditions deviating from the ideal situation [1], the algorithm is strictly valid only if the following experimental conditions are fulfilled:

1. The object imaged consists of a single, homogeneous material.
2. Monochromatic radiation is used.
3. The distance z between the object and the detector plane fulfils the near-field condition, i.e.,

$$z \ll d^2/\lambda,$$

where z is the distance between object and detector, d is the characteristic size of the smallest discernible features in the object, and λ is the X-ray wavelength.

If these conditions are met, the intensity distribution $I(x, y)$ measured at a single, known distance z between the object and the detector plane can then be used to retrieve the projected thickness $t(x, y)$ of the object (or, which is equivalent, the projected phase shift of the X-ray wavefront) using the relation

$$t(x, y) = -\frac{1}{\mu} \ln \left(\mathcal{F}^{-1} \left[\frac{\mu \mathcal{F}[I(x, y)/I_0(x, y)]}{\mu + z\delta(u^2 + v^2)} \right] \right).$$

Here, x and y are the Cartesian coordinates in the image and/or object plane, μ is the linear attenuation coefficient of the object material, \mathcal{F} and \mathcal{F}^{-1} are, respectively, the forward and backward Fourier transform operators, $I(x, y)$ is the intensity distribution in the phase-contrast

radiograph, $I_0(x, y)$ is the incident intensity just upstream of the object, δ is the decrement from unity of the X-ray refractive index of the object material, u, v are the complex conjugate coordinates of x and y , and z is the object–detector distance. The quantities δ and μ are related to the complex-valued X-ray refractive index n by

$$n = 1 - \delta + i\beta$$

and

$$\mu = 4\pi\beta / \lambda ,$$

where both β and δ are dimensionless real numbers, and λ is the X-ray wavelength.

1.3 Flat-field and dark-image correction

In a real X-ray imaging system, the background signal (i.e., signal recorded if no object is in the beam) is generally not uniform. This can be due to inhomogeneities in the X-ray beam intensity distribution or to non-uniform detector response (for example due to vignetting or varying pixel sensitivities), or both. Moreover, the quantity of interest for any quantitative analysis of the radiographs is generally the relative transmitted intensity, I/I_0 .

It is therefore generally required to normalize each radiograph by an image of the beam without the object. This is called a *flat-field correction*. Also, in each image recorded with the detector, the offset signal recorded even with no photons hitting the detector (*dark signal*) should be subtracted from the image before further processing (*dark-image correction*). The corrected images are

$$I_c(x, y) = \frac{I(x, y)}{I_0(x, y)} = \frac{S_o(x, y) - S_d(x, y)}{S_f(x, y) - S_d(x, y)} ,$$

where $S_o(x, y)$ is the image signal measured with the object in the beam, $S_d(x, y)$ is the dark signal, and $S_f(x, y)$ the flat signal (i.e., the image with the sample removed).

In order to keep the increase of statistical noise by the correction to a minimum, it is preferable to take a series of flat-field and dark images (rather than just one of each) and average these before using them for correcting the object images.

ANKAphase can perform flat-field and dark correction of the radiographs. It can process series of flat-field images, as well as series of dark images by calculating, for every pixel, the average gray value for that pixel, where the averaging is done over all images in the series (but separately for each pixel). In doing this, ANKAphase uses median averaging (rather than the arithmetic mean).

1.4 Boundary treatment; image padding

To avoid crosstalk between opposed ends of the images, the flat-corrected projection images $I_c(x, y)$ are padded to larger size before applying the phase retrieval algorithm. Without padding, crosstalk occurs because the convolution of the projection data with a phase-reconstruction kernel is carried out in Fourier space (via the reconstruction formula in section 1.2), so that the convolution is periodic. Further padding is required because the fast Fourier transform (FFT)

routine used by *ANKAphase* requires the number of pixels in each dimension to be a power of two.

In the numerical implementation of the phase-reconstruction algorithm, the corrected projection radiographs containing the normalized intensity distribution $I_c(x, y)$ are therefore extended to larger size by a surrounding margin before the reconstruction formula from section 1.2 is applied. Thus, each of the two dimensions of the image is padded from its original value of n_{original} pixels to a new width n_{padded} of

$$n_{\text{padded}} = \exp\{(\ln 2) \times \text{ceil}[\log_2(n_{\text{original}} + n_{\text{ext}})]\} ,$$

where

$$n_{\text{ext}} = 2 \times \text{ceil} \left[\frac{3\lambda z}{(\Delta x)^2} \right] .$$

Here, $\text{ceil}(x)$ is the ceiling function, which yields the smallest integer number that is equal to or greater than x . The X-ray wavelength and propagation distance are denoted, respectively, λ and z , as before, and Δx is the pixel size of the detector.

Each flat-corrected projection radiograph is thus first extended by margins of width $n_{\text{ext}}/2$ on all four sides and then padded further up so that its dimensions are powers of two.

Each pixel in the padded area is filled with the value of the nearest pixel in the original image.

The formula for n_{ext} takes into account the radius of the area around each position in the detector plane from which diffraction effects by features in the object plane are not negligible. A derivation of the formula can be found elsewhere [1].

A checkbox option allows the user to force $n_{\text{ext}} = 0$ (“no automatic padding”, see section 5.3.3). If this option is set, the images will still be padded to a size corresponding to a power of two. This option is included for cases in which the images have already been padded before processing with *ANKAphase*. It should not be used on images to which no boundary treatment has been applied before.

In any case, before output, the reconstructed images are cropped back to the original dimensions of the input images.

2 Where to get *ANKAphase*

ANKAphase can be downloaded from the ImageJ plugin web site hosted by NIH. The address is <http://rsbweb.nih.gov/ij/plugins/>.

3 Files in the distribution

The *ANKAphase* software distribution package contains the following files. We recommend that you place all of these files in the same directory.

ANKAphase_2_0.jar The Java application executable.

runANKAphase.bat A batch file recommended to launch *ANKAphase* as a standalone file on Windows systems.

runANKAphase.command A shell script recommended to launch *ANKAphase* as a standalone file on Mac OS X systems.

runANKAphase.sh A shell script recommended to launch *ANKAphase* as a standalone file on GNU/Linux or other Unix systems.

Note: You can edit the launcher batch file or shell script with a standard text editor to set paths according to your needs, or to decide for yourself how much heap space Java will allocate for *ANKAphase* when it is run as a standalone application.

ankaphase-jsr2011-preprint.pdf A preprint of a peer-reviewed article [1] on *ANKAphase* (can be viewed and/or printed with Adobe Reader, xpdf, Foxit Reader or other PDF viewer software).

ankaphase-userguide.pdf The User Guide, as a Portable Document Format file (can be viewed and/or printed with Adobe Reader, xpdf, Foxit Reader or other PDF viewer software).

ankaphase-userguide.html The User Guide, in Hypertext Markup Language, for viewing in Web browsers such as Firefox, Opera, Safari etc.

ankaphase-userguide-e0[1-5].gif Five small auxiliary graphics files used by the HTML User Guide file.

4 Requirements

4.1 Platforms and operating systems

ANKAphase being written in Java, it should run on any platform and with any operating system for which Java is available, i.e., Microsoft Windows, Mac OS X, GNU/Linux, and Sun Solaris systems. However, the program has so far been tested only on Microsoft Windows XP, GNU/Linux and Mac OS X (Snow Leopard) systems.

4.2 Software requirements

In order to run *ANKAphase*, you need a **Java Runtime Environment** (JRE) version 6 or higher (Java version 1.6.0 or higher) [3], installed on your system. There is a high probability that Java and the JRE are already installed on your computer. If they are not, you can download them free of charge from Sun Microsystems [4].

You can use *ANKAphase* as a **plugin for ImageJ**. If you want to do this, you will need ImageJ installed on your system. Like *ANKAphase*, ImageJ is programmed in Java. It is open-source software and can be downloaded from the NIH web site [5].

4.3 Data format requirements

The input data should be stacks of two-dimensional image files, containing a single image frame per file. *ANKAphase* can read and write the following image file formats (when reading images, it determines the format from the file-name extension):

- TIFF (unsigned integers with 8 or 16 bit depth, or single-precision float data),
- JPEG (8-bit unsigned integer),
- PNG (8-bit unsigned integer),
- BMP (8-bit unsigned integer),
- and the ESRF-specific EDF format (4-byte single-precision float data).

5 Getting started

5.1 Running it as a standalone application

Used as a standalone application, *ANKAphase* does not need any particular installation at all. It runs perfectly well if you don't "install" it. Just make sure that the files in the distribution — preferably all of them (see section 3 above), but at least the Java application executable (*ANKAphase_2_0.jar*) and the launcher batch file (*runANKAphase.bat* for Windows, *runANKAphase.command* for Mac OS X, *runANKAphase.sh* for GNU/Linux or other Unix-like systems) — are in the same directory. Then open that directory and start the launcher batch file (typically by double-clicking on its name, or by entering its name at the command prompt).

If you want the "Help" function in *ANKAphase* to work properly, the files whose names start with *ankaphase-userguide* should also be placed in the same directory as *ANKAphase_2_0.jar*.

What the batch file does: Actually, *ANKAphase* will start up fine when launched without the batch file (by just opening *ANKAphase_2_0.jar* directly). However, this will usually cause Java to allocate only a very limited amount of "heap space" (memory used for data). What the batch file does is cause Java to allocate a larger amount of heap space than it does by default. This ensures that *ANKAphase* will work properly when used with large amounts of data. We therefore strongly recommend using the batch file when launching *ANKAphase* as a standalone application.

Note: You can edit the batch file with a standard text editor to set paths according to your needs, or to decide for yourself how much heap space Java will allocate for *ANKAphase* when it is run as a standalone application.

5.2 Running it as a plugin for ImageJ

Place the *ANKAphase_2_0.jar* file in the subdirectory *plugins* of your ImageJ installation directory. (On a Microsoft Windows system, this could be "C:\Program Files\ImageJ\plugins".)

When running ImageJ, an entry named “*ANKAphase*” will then be available in the “Plugins” menu or in one of its submenus.

If you want the “Help” function in *ANKAphase* to work properly, the files whose names start with *ankaphase-userguide* should be placed in the same directory as *ANKAphase_2_0.jar*.

5.3 Using it

All of *ANKAphase*’s control elements are contained in a single control window that will appear when you launch the application. The window is thematically divided into three areas: “General settings”, “Flat-field settings”, and “Phase-retrieval settings”. Some additional settings can be made near the bottom of the window, where you will also find the “Run” and “Stop” buttons. There is also a menu bar at the top of the window. The following sections explain these control elements in more detail.

5.3.1 General settings

The “General settings” area in the upper part of the *ANKAphase* control window is where you

- set the directory paths in which *ANKAphase* will look for its input data,
- choose whether or not to perform dark-image subtraction, and whether and how to carry out flat-field correction.

How *ANKAphase* finds its input images *ANKAphase* identifies the input image files it will process by the directory in which these files are located. For a given set of images taken under the same experimental conditions and with the same image dimensions for all frames, it expects all the projection radiographs to be in one directory, while the series of dark images to be averaged for the dark correction are supposed to be in another directory.

For the flat-field images, *ANKAphase* foresees two directories: one in which flat images taken before the acquisition of the projections are stored, and another in which flats taken after the acquisition of the projections are stored. If two blocks of flat-field images obtained in this manner are available, *ANKAphase* will correct the first half of the projection images using the first set of flat-field images, and the second half of the projections will be corrected using the second set of flat-field images (unless the “Interpolate flat-field images”) option is set – see section 5.3.2 below).

It is the user’s responsibility to ensure that the dimensions and data type of all of the images in the input directories are the same.

Projection images directory This is the directory in which *ANKAphase* looks for the projection radiographs. The program will scan this directory for all files whose filename extension matches one of the image formats allowed for input. All of these images are then processed.

Dark-field images directory This is the directory in which *ANKAphase* looks for dark images. All images found in this directory are averaged (pixel-wise, median average) to form an average dark image that will then be used for the dark correction.

The checkbox on the left of this field controls whether a dark correction will be performed or not. Dark correction will only be made if the box is checked.

Flat-field 1 images directory This is the directory in which *ANKAphase* looks for flat-field images taken *before* the acquisition of the projection radiographs to be processed. The flat images are averaged in the same way as the dark image series, and the averaged image is then used for the correction.

The checkbox on the left of this field controls whether the images in the “Flat-field 1 images directory” will actually be read, processed and used for flat-field correction or not.

Flat-field 2 images directory This is the directory in which *ANKAphase* looks for flat-field images taken after the acquisition of the projection radiographs to be processed. The images are averaged in the same way as the dark image series and the first block of flat-field images, and the averaged image is then used for the correction.

The checkbox on the left of this field controls whether the images in the “Flat-field 2 images directory” will actually be read, processed and used for flat-field correction or not.

Note: If none of the checkboxes on the left of “Flat-field 1 images directory” and “Flat-field 2 images directory” is checked, no flat-field correction will be made. If only one of the two checkboxes is checked, all projection images will be flat-field-corrected using the images in the flat-field directory checked. If both boxes are checked and the option “Interpolate flat-field images” (see section “Flat-field settings” below) is not selected, the first half of the projection image files will be corrected using the “flat-field 1” images, and the second half using the “flat-field 2” images. If both boxes are checked and the option “Interpolate flat-field images” is selected, each projection image will be corrected by a linear combination of the averaged flat-field images from each of the two flat-image directories.

5.3.2 Flat-field settings

The “Flat-field settings” area of the control window is where you

- choose whether or not to save flat-field corrected images to disk, and,
- if flat-field correction is performed (see section 5.3.1 for how to control this), whether or not the correction should be made using interpolated flat-field images.

Save flat-corrected images Check this box if you want to write out files containing the projection images after flat-field and dark correction (but prior to phase reconstruction). The following parameters can only be set if this option is selected:

Output directory The path of the directory to which the corrected projection files will be written.

File prefix The filename prefix of the corrected projection files. The entire file name of each output file will consist of this prefix, an underscore sign, a four-digit number, and the filename extension corresponding to the file format. The numbering starts at the number indicated in the field “Output file start number” near the bottom of the window. (The numbering of the input files is ignored in the output file naming.)

For example, if you enter “out” in this field, and have chosen TIFF as the output format, the reconstruction files will have names *out_0000.tiff*, *out_0001.tiff*, and so forth.

Format The image file format, data type and bit depth for the output files. “Float”-type images are stored in single precision, i.e., with a depth of 4 bytes per pixel.

Scaling option If you have chosen an output format with integer data type (i.e., any except the “float” formats), this option list allows you to select how the values of the corrected projections, which are internally calculated as floating-point values, are mapped to the range of values allowed by the output data type.

Be aware that the scaling operation, while necessary for integer-type output, will usually result in a loss of information due to the reduced dynamic range of the output data type.

In order for you to be able to retrieve the approximate original floating-point values from the images afterwards, a text file named *Scaling_flat_images.txt* will be written to the output directory. This file contains the original minimum and maximum values of each output image, before the conversion from floating-point to integer data type.

Caveat: The file *Scaling_flat_images.txt* will be overwritten without backup upon every new run of ANKAphase, unless you have changed the output directory.

The possible scaling options are:

Scale each image to its min/max This is the default behavior. The values of each image will be scaled so the the full range of values in the integer output image is used. This ensures minimum loss of information, although it means that the scaling parameters will usually be different for each image in the series of images processed. Use the information in the file *Scaling_flat_images.txt* (see above) if you need to retrieve the quantitative values of the corrected images later.

Scale to $3 \times (\text{max-min})$ of first image This will use a global set of scaling parameters for all images in the series. The parameters are defined automatically so that the output image range is three times the range of the first processed image in the series. The assumption behind this is that the original range of values in the series of images may change, but the change is not excessive. As with the other scaling options, the information in the file *Scaling_flat_images.txt* (see above) can be used if you need to retrieve the quantitative values of the corrected images later.

Scale to user-specified value range This will use a global set of scaling parameters defined by the user via the values entered in the “from:” and “to:” fields. Set these two values to the expected minimum and maximum floating-point values of all corrected projections. As with the other scaling options, the infor-

mation in the file *Scaling_flat_images.txt* (see above) can be used if you need to retrieve the quantitative values of the corrected images later.

Interpolate flat-field images If this option is set, each projection image is corrected with a flat-field image calculated by pixel-wise linear interpolation of the averaged flat-field image from the first block and the averaged flat-field image from the second block of flat-field images. If the option is not set (the default setting), the first half of the projection images will be corrected using the first set of flat-field images, and the second half of the projections will be corrected using the second set of flat-field images.

5.3.3 Phase-retrieval settings

The “Phase-retrieval settings” area of the control window is where you

- choose whether or not phase reconstruction should be carried out,
- enter the physical parameters that will be used for phase reconstruction,
- select the output file path.

Save phase-corrected images Check this box if phase retrieval should be carried out and the phase-retrieved projection files written out to files. (Not selecting this option can make sense if you want to use *ANKAphase* to only perform flat-field and dark-image correction on a series of projection images.)

The parameters that you need to set if this option is selected are:

Experiment parameters: The physical parameters of the experimental setup and X-ray optical properties of the sample.

beta: The imaginary part β of the estimated complex-valued X-ray index of refraction (see section 1.2 above) of the sample material, in units of 10^{-9} . For example, if $\beta = 1.73 \times 10^{-8}$, you should enter “17.3” here.

delta: The real-part decrement δ of the estimated complex-valued X-ray index of refraction (see section 1.2 above) of the sample material, in units of 10^{-6} . For example, if $\delta = 8.95 \times 10^{-7}$, you should enter “0.895” here.

Note: If you don’t know the X-ray refractive index of the material your sample consists of, but you do know (or have an idea of) the elemental composition and density of the material, you can look up the quantities δ and β in tabulated databases, some of which can be accessed or downloaded online and free of charge [6, 7].

distance: The distance between the sample and the detector plane, in millimeters (mm).

energy: The mean effective X-ray photon energy, in kiloelectronvolts (keV).

pixel size: The effective pixel size of the detector, in microns (μm).

Output directory The path of the directory to which the corrected projection files will be written.

File prefix, Format, Scaling option The filename prefix of the phase-retrieved projection files, their image data format and data type, and the method of value rescaling. For details, see the description of the corresponding fields in the “Flat-field settings” (section 5.3.2 above). Note that the output text file containing the scaling parameters applied (*Scaling_flat_images.txt* for the flat images) is named *Scaling_phase_images.txt* for the phase-retrieved images.

Caveat: The file *Scaling_phase_images.txt* will be overwritten without backup upon every new run of ANKAphase, unless you have changed the output directory.

No automatic padding If this option is set, the images will not be automatically extended by a margin of width adapted to the experimental parameters prior to application of the phase retrieval transform. See section 1.4 above for more details.

Note: You should only use this option if you have applied your own adapted edge extension to the projection images before processing with ANKAphase.

5.3.4 Other settings

The bottom region of the control window allows you to set a few more options:

Calculate images from <n1> to <n2> : Select this option to define a limited range of input files to be processed. If this option is not selected, all image files in the “projection images” directory will be processed. If it is selected, only the <n1>st to the <n2>st projection image will be processed. Counting starts at 1.

Show images If this option is selected, the processed images will be displayed on the screen during processing. Note that images displayed may appear black on screen when a “float”-type output file format is selected. If this happens, it does not mean that the image files are not okay.

CPUs Select the maximum number of CPUs used in parallel by ANKAphase.

Output file start number The file number offset of the output image files (see also the explanation on output file naming and numbering above, section 5.3.2). This can be handy in cases where more than two blocks of flat-field images exist in the primary data. Such cases are not handled automatically by ANKAphase, but you can use the “Calculate images from ... to ...” option together with the “Output file start number” option to perform flat-field correction in the desired way, i.e., in batches of images.

5.3.5 Launching the processing

Once all parameters are set, use the “Run” button on the lower right of the control window to launch data processing and writing to the output files.

The status bar at the bottom of the control window will show the progress of the processing. If the “Show images” option has been selected, display windows with the processed images will pop up during processing. After processing has finished, a small alert window will pop up informing you how many output image files have been written to disk.

The “Stop” button to the left of the “Run” button allows you to abort processing.

5.3.6 Menu bar

The menu bar (at the top of the ANKAphase window) has only two entries: “File” and “Help”. While none of them is strictly necessary to use the program, both give access to functionality that may be very useful.

“File” menu The items in the “File” drop-down menu allows you to

- save a set of parameters for later use (“Save parameters”),
- load a parameter set previously saved (“Load parameters”), or
- quit *ANKAphase* (“Exit”).

Note: **Parameter files** are written as Ascii files and store the current contents of all the control elements in the *ANKAphase* window. When *ANKAphase* starts, it will look for a parameter file named *parameter.txt* in the directory from which it was started (usually the same directory in which the .jar executable is located). If *parameter.txt* exists, it will automatically be loaded. When saving a parameter file through the “Save parameters” menu entry, you can give it any name. If you name it *parameter.txt* and store it in the *ANKAphase* distribution directory, it will be loaded automatically the next time *ANKAphase* is started.

Note that *ANKAphase* saves its parameters automatically when exiting, but only when processing was launched at least once.

Note also that saving the parameters (automatically or via the “Save parameters” entry in the File menu) will overwrite any potentially existing parameter files of the same name without asking.

“Help” menu The “Help” drop-down menu allows you to

- display some basic information about the program (“About”),
- open the online help file for *ANKAphase* (i.e., the text you are reading) as an HTML file in the standard web browser of your system (“Help”). In order for this feature to work properly, you should place all the help files in the distribution (files whose names start with *ankaphase-userguide*) in the same folder as *ANKAphase_2_0.jar*.

6 Further information; referencing

Beyond this User Guide, an article in the Journal of Synchrotron Radiation [1] further describes the details of the algorithm used here, the conditions under which it yields good results, and gives application examples. If you find *ANKAphase* useful and present or publish results obtained with it, we kindly ask you to give us proper credits by citing the journal article.

References

- [1] T. Weitkamp, D. Haas, D. Wegrzynek, and A. Rack, *ANKAphase: software for single-distance phase-retrieval from inline X-ray phase contrast radiographs*, Journal of Synchrotron Radiation (2011), vol. 18, issue 4, p. 617-629, DOI: 10.1107/S0909049511002895.
- [2] D. Paganin, S. C. Mayo, T. E. Gureyev, P. R. Miller, and S. W. Wilkins, *Simultaneous phase and amplitude extraction from a single defocused image of a homogeneous object*, Journal of Microscopy **206** (2002) 33–40, DOI: 10.1046/j.1365-2818.2002.01010.x.
- [3] If you are unsure which version of Java is installed on your system, you can find out by typing “java -version” at the command prompt of your system (this works both on Windows [on Windows XP, the command prompt is accessible via “All Programs – Accessories – Command prompt”] and on Unix systems).
- [4] <http://www.java.com/download/>, the download address for the Java Runtime Environment (JRE).
- [5] <http://rsb.info.nih.gov/ij/>, the ImageJ web site.
- [6] http://henke.lbl.gov/optical_constants/, an online calculator of X-ray optical properties at the Center for X-Ray Optics (CXRO) at Lawrence Berkeley Laboratory in Berkeley, California, USA.
- [7] <http://www.esrf.eu/UsersAndScience/Experiments/TBS/SciSoft/xop2.3>, the web site of XOP (X-Ray Oriented Programs), a software suite for X-ray optical calculations including a database of X-ray optical properties (DABAX). XOP is maintained by two synchrotron light laboratories, the European Synchrotron Radiation Facility (ESRF) in Grenoble, France and the Advanced Photon Source (APS) in Argonne, Illinois, USA.